In the claims:

Following is a complete set of claims as amended with this Response.

1. (Previously Presented) A method comprising:

receiving packets from an external network;

determining a destination for each packet through the switch fabric;

determining a path to be taken by each packet through a switch fabric;

classifying each packet into one of a plurality of flow bundles based on the packet's destination and path through the switch fabric and labeling each packet with a flow identifier to identify the associated flow bundle;

mapping each classified and labeled packet into one of a plurality of queues;

queuing each mapped packet into the respective mapped queue to await transmission based on the flow bundle into which the packet has been classified so that all packets in a queue belong to a same flow bundle;

regulating rates at which packet traffic moves out of the queues with a traffic shaping algorithm; and

scheduling the packets in the queues for transmission from the queues to a next destination through the switch fabric.

2. (Cancelled)

3. (Original) The method of claim 1, wherein determining a path to be taken by each packet through a switch fabric comprises determining a path to be taken by each packet through a switch fabric based on load balancing.

4. (Currently Amended) The method of claim 1, wherein classifying each packet further comprises ~~comprising~~ labeling each packet with information identifying an associated flow .

5. (Original) The method of claim 1, wherein classifying each packet into one of a plurality of flow bundles comprises classifying each packet into one of a plurality of flow bundles based on the packet's destination, path through the switch fabric, and priority.

6. (Original) The method of claim 1, wherein scheduling the packets in the queues for transmission comprises scheduling the packets in the queues for transmission using a Round Robin scheduling algorithm.

7. (Original) The method of claim 1, wherein scheduling the packets in the queues for transmission comprises scheduling the packets in the queues for transmission using a Longest Delay First algorithm.

8. (Original) The method of claim 1, wherein scheduling the packets in the queues for transmission comprises scheduling the packets in the queues for transmission using a Stepwise QoS Scheduler (SQS).

9. (Previously Presented) The method of claim 1, further comprising determining a traffic class to which each received network packet belongs based on protocols associated with the packet.

10. (Currently Amended) The method of claim 1, wherein scheduling the packets in the queues further comprises ~~comprising~~ forwarding the packets to a switch coupled to the switch fabric for transmission to the next destination.

11. (Previously Presented ) An apparatus comprising:

an input to receive packets from an external network;

a classification unit to examine the received packets, determine a destination for each packet through the switch fabric, determine a path to be taken by each packet through a switch fabric, classify each packet into one of a plurality of flow bundles based on the packet's destination and path through the switch fabric, and label each packet with a flow identifier to identify the associated flow bundle;

a mapping unit coupled to the classification unit to place each classified and labeled packet into one of a plurality of queues based on the flow bundle into which the packet has been classified so that all packets in a queue belong to a same flow bundle;

one or more traffic shapers coupled to the mapping unit to regulate the rate at which packet traffic moves out of the queues; and

a scheduler coupled to the traffic shapers to regulate the order in which packets in the queues will be transmitted from the queues to a next destination through the switch fabric.

12. (Original) The apparatus of claim 11, further comprising an access unit coupled to the classification unit to receive packets from and transmit packets to the network.

13. (Original) The apparatus of claim 11, further comprising a switch coupled to the scheduler to transmit the scheduled packets to the switch fabric.

14. (Original) The apparatus of claim 11, wherein the classification unit comprises a load balancing element to determine a path to be taken by each packet through a switch fabric based on load balancing.

15. (Original) The apparatus of claim 11, wherein the classification unit comprises a labeling element to label each packet with information identifying an associated flow and flow bundle.

16. (Previously Presented) A network processor microengine of a network switching node for running a plurality of threads to perform processes on received packets, the processes of the threads comprising:

determining a path to be taken by each received network packet through a switch fabric;

classifying each packet into one of a plurality of flow bundles based on the packet's destination and path through the switch fabric;

labeling each packet with a flow identifier to identify the associated flow bundle;

mapping each packet into one of a plurality of queues;

queuing each packet into the respective mapped queue to await transmission based on the flow bundle into which the packet has been classified;

regulating a rate at which traffic moves out of the queues using a traffic shaping algorithm; and

scheduling the packets in the queues for transmission from the queues to a next destination through the switch fabric.

17. (Cancelled)

18. (Previously Presented) The network processor microengine of claim 16, wherein the process of the threads further include labeling each packet with information identifying an associated flow and flow bundle.

19. (Previously Presented) The network processor microengine of claim 16, wherein the process of the threads further include determining a traffic class to which each received network packet belongs.

20. (Previously Presented) The network processor microengine of claim 16, wherein the process of determining a path to be taken by each received network packet through a switch fabric comprises determining a path to be taken by each received network packet through a switch fabric based on load balancing.

21. (Previously Presented) The network processor microengine of claim 16, wherein the process of classifying each packet into one of a plurality of flow bundles comprises classifying each packet into one of a plurality of flow bundles based on the packet's destination, path through the switch fabric, and priority.

22. (Previously Presented) The network processor microengine of claim 16, wherein the the process of the threads further include forwarding the packets to a switch coupled to the switch fabric for transmission to the next destination.

23. (Previously Presented) A system comprising:

a switch to receive from an external network and transmit packets through a switch fabric to a destination;

a classification unit to examine packets received from a network through the switch, determine a path to be taken by each packet through the switch fabric, and classify each packet into one of a plurality of flow bundles based on the packet's destination and path through the switch fabric;

a mapping unit coupled to the classification unit to place each classified packet into one of a plurality of queues based on the flow bundle into which the packet has been classified so that all packets in a queue belong to a same flow bundle;

one or more traffic shapers coupled to the scheduler to regulate rates at which traffic moves out of the queues;

a scheduler coupled to the mapping unit to regulate the order in which packets in the queues will be transmitted from the queues to a next destination; and

24. (Cancelled)

25. (Original) The system of claim 23, wherein the classification unit comprises a load balancing element to determine a path to be taken by each packet through the switch fabric based on load balancing

26. (Original) The system of claim 23, wherein the classification unit comprises a labeling element to label each packet with information identifying an associated flow and flow bundle.

27. (Previously Presented) The method of claim 1, wherein classifying further comprises classifying into one of a plurality of traffic classes based on a type of traffic.

28. (Previously Presented) The method of claim 6, wherein queuing each mapped packet comprises queuing so that all packets in a queue have the same priority.

29. (Previously Presented) The apparatus of claim 11, wherein the classification unit further classifies packets into one of a plurality of traffic classes based on a type of traffic.

30. (Previously Presented) The apparatus of claim 11, wherein the classification unit determines a priority for each packet and wherein the mapping unit queues each mapped packet so that all packet in a queue have the same priority.